server processes a first user request and a second server processes a second user request. The servers are connected to a switch, which is in turn connected to one or more data storage devices.

In one aspect of the invention, the system is fully "scalable" in the sense that additional servers can be added to the system as demand for a particular application increases, without adding additional data storage devices. Conversely, servers can be removed from the system without removing data storage devices. In a similar manner, additional data storage devices can be added to the system as storage requirements for a particular application increase, without adding additional servers. Conversely, data storage devices can be removed from the system without removing servers. Thus, the system is scalable to increase or decrease server capacity without changing the data storage capacity, and/or is scalable to increase or decrease data storage capacity without changing the server capacity.

In another aspect of the invention, each of the servers applies substantially the same application(s) when processing user requests, and each of the data storage devices contains substantially identical data. The system is fully "survivable" in the sense that, if any one of the servers fails, user requests can be processed by any of the other servers in the system that are operable. Likewise, if any one of the data storage devices fails, substantially identical data can be retrieved from any of the other data storage devices that are operable. Thus, the system is survivable and able to process user requests in the event of a failure of any one of the servers, and is survivable and able to retrieve data in the event of a failure of any one of the data storage devices.

Rejections Under 35 U.S.C. § 102

In the Office Action dated April 8, 2002, the Examiner rejected claims 79-80, 82-88, 91-95, 98-99, 103-104, 106 and 108-114 under 35 U.S.C. § 102 (e) as being anticipated by

U.S. Patent No. 5,699, 503 to Bolosky et al. ("Bolosky"). Bolosky discloses a media server system, such as a video-on-demand system in which video image sequences (*e.g.*, a movie) are transmitted from the system to subscribers in response to user requests. The system includes a controller (*e.g.*, controller 16 of Fig. 2) and a plurality of subsystems (*e.g.*, subsystems 18A, 18B and 18C of Fig. 2). Bolosky, col. 5, l. 61 to col. 6, l. 10. Each subsystem comprises a single microprocessor (*e.g.*, microprocessor 20A of Fig. 2) and one or more data storage devices (*e.g.*, data storage devices 22A and 24A of Fig. 2). Bolosky, col. 6, ll. 11-19. In operation, the controller cooperates with the microprocessor of each of the subsystems to schedule the transmission of video image sequences stored on the data storage devices to the subscribers. Bolosky, col. 6, ll. 20-23.

The video image sequences are stored on the data storage devices by dividing them into sequential blocks of data and "striping" them across the primary portions of the data storage devices. Bolosky, col. 6, ll. 40-43. "Striping" refers to the method in which a first block of data is stored on a first data storage device and each sequentially following block of data is stored on the next sequential data storage device. Bolosky, col. 6, ll. 46-49. When reaching the last data storage device, the next block of data wraps around and is stored on the first data storage device. Bolosky, col. 6, ll. 49-51. This continues until all the blocks of data are stored across the data storage devices. Bolosky, col. 6, ll. 51-53.

After the blocks of data are stored on the primary portions of the data storage devices, "declustered mirroring" is used to store a copy of this same data on the secondary portions of the data storage devices. Bolosky, col. 6, ll. 57-61. An example of "declustered mirroring" is shown in Fig. 3D. In this example, a block of data stored on the primary portion of a data storage device of a first subsystem (*e.g.*, "Subsystem 1, SD1, Block A") is divided into

first and second sub-blocks of data, wherein the first sub-block of data is stored on the secondary portion of a data storage device of a second subsystem (*e.g.*, "Subsystem 2, SD3, Sub-Block A1") and the second sub-block of data is stored on the secondary portion of a data storage device of a third subsystem (*e.g.*, "Subsystem 3, SD5, Sub-Block A2"). Bolosky, col. 9, lines 14-34. In the event of a failure of the first subsystem, the first and second sub-blocks of data can be transmitted from the second and third subsystems to the subscribers. Id.

Independent claims 79, 82, 91, 102, 104 and 106 (and dependent claims 80-81, 83-89, 92-97, 103, 105 and 107-113) of the present invention, which are directed to the "scalability" aspect of the invention,[1] each include the limitation that the server(s) operate independently of the data storage device(s) so as to permit the addition (or removal) of a server without the addition (or removal) of a data storage device (*e.g.*, as demand for a particular application increases or decreases). Bolosky does not disclose or suggest this limitation.

The Examiner cites to three different portions of the Bolosky specification to support his contention that Bolosky discloses a system that is fully scaleable. However, none of these portions discloses the addition (or removal) of a server without the addition (or removal) of a data storage (*e.g.*, as demand for a particular application increases or decreases):

1. Bolosky, col. 8, l. 38 to col. 9, l. 34

This portion of the Bolosky specification discloses that a declustering number (*i.e.*, the number of sub-blocks of data stored across multiple data storage devices) can be chosen so as to tolerate the failure of more than one data storage device or subsystem. An alternative embodiment of "declustered mirroring" is also

[1] Dependent claims 86-88, 103, 109, 111 and 113 are also directed to the "survivability" aspect of the invention.

disclosed, wherein the burden of performing failure mode processing is spread across a larger number of data storage devices than in the preferred embodiment.

2. Bolosky, col. 5, l. 61 to col. 6, l. 23

This portion of the Bolosky specification discloses that, although the preferred embodiment describes three subsystems, a larger number of subsystems will typically be employed. This portion also discloses that while each subsystem of the preferred embodiment includes a single microprocessor that is responsible for controlling two data storage devices, each subsystem may alternatively include one data storage device or more than two data storage devices.

3. Bolosky, col. 7, ll. 4-28

This portion of the Bolosky specification discloses that the declustering number may vary, and that a higher declustering number can be chosen to: (1) lessen the burden of failure mode processing by any one data storage device, and (2) reduce the bandwidth reserved for failure mode processing.

All of the above portions of the Bolosky specification are directed to the system's capability of varying the number of data storage devices within a subsystem to enhance fault tolerance. Nowhere does Bolosky disclose the addition (or removal) of a server (*i.e.*, microprocessor) without the addition (or removal) of a data storage device (*e.g.*, as demand for a particular application increases or decreases). Rather, in Bolosky, each microprocessor is tied to particular data storage devices (albeit the number of data storage devices may vary) and does not operate independently of those data storage devices. Thus, because Bolosky does not disclose or suggest this limitation, these claims are patentable over Bolosky.

Independent claims 98 and 114 (and dependent claims 86-88, 99-101, 103, 109, 111, 113 and 115) of the present invention, which are directed to the "survivability" aspect of the invention,[2] each include the limitation that each of the data storage devices stores substantially the same data such that, in the event of a failure of any one of the data storage devices, the data is accessible from any other of the data storage devices that are operable. Bolosky does not disclose or suggest this limitation.

The Examiner argues that the "declustered mirroring" process of the Bolosky system discloses this limitation. Not true. In Bolosky, each of the data storage devices do not store substantially the same data. In fact, as can be seen in Fig. 3D, each data storage device stores a completely different set of data (SD1 stores Block A and Sub-Blocks I1 and G2; SD2 stores Block B and Sub-Blocks F1 and D2; etc.). This configuration, in which different blocks of data are stored on each data storage device for survivability, is directly contrary to the claimed configuration of storing substantially the same data on each of the data storage devices. Therefore, these claims are also patentable over Bolosky.

Rejections Under 35 U.S.C. § 103

The Examiner rejected claims 81, 89, 96, 100, 105, 107 and 115 under 35 U.S.C. §103(a) as being obvious over Bolosky. The Examiner also rejected claims 90, 97, 99 and 101-103 under 35 U.S.C. § 103(a) as being obvious over Bolosky in view of U.S. Patent No. 4,914,570 to Peacock ("Peacock"). However, for the reasons stated above with respect to the § 102 rejection, these claims are distinguishable from Bolosky and/or Peacock.

---

[2] Dependent claims 86-88, 103, 109, 111 and 113 are also directed to the "scalability" aspect of the invention.

In view of the foregoing remarks, it is respectfully submitted that the claims are now in condition for allowance and eventual issuance. Such action is respectfully requested. Should the Examiner have any further questions or comments which need be addressed in order to obtain allowance, he is invited to contact the undersigned attorney at the number listed below.

Respectfully submitted,

By: Judith L. Carlson

Judith L. Carlson, Reg. No. 41,904
STINSON MORRISON HECKER LLP
1201 Walnut Street, Suite 2800
P.O. Box 419251
Kansas City, MO 64141-6251
Telephone: (816) 842-8600
Facsimile: (816) 691-3495
Attorney for Assignee